arXiv:1302.1634v1 [math.CO] 7 Feb 2013

HYPERGRAPH LIMITS VIA MARTINGALES

YUFEI ZHAO

Abstract. A sequence of k-uniform hypergraphs $H_1, H_2, \dots$ is said to converge if the homomorphism densities $t(F, H_1), t(F, H_2), \dots$ converge for every k-uniform hypergraph F. Lovász and Szegedy showed that every convergent sequence of graphs has a limit in the form of a symmetric measurable function $W\colon [0,1]^2 \to [0,1]$ that determines all limits of subgraph densities. For hypergraphs, analogous limits $W\colon [0,1]^{2^k-2} \to [0,1]$ were constructed by Elek and Szegedy using ultraproducts.

We give a new ultraproduct-free construction of hypergraph limits as well as a colored generalization. Our proof follows the original approach of Lovász and Szegedy, with the key ingredients being a weak Frieze-Kannan type regularity lemma and the Martingale Convergence Theorem.

1. Introduction

One of the starting points in the theory of dense graph limits is the seminal paper by Lovász and Szegedy [9] where they constructed limit objects for convergent sequences of dense graphs. The subject has grown enormously since then with many exciting developments (see Lovász's recent monograph [7]).

For any two graphs F and G, let $\hom(F, G)$ denote the number of homomorphism from F to G, i.e., maps $V(F) \to V(G)$ that carry every edge of F to an edge of G. The homomorphism density $t(F, G)$ is defined to be the probability that a random $V(F) \to V(G)$ is a homomorphism, i.e.,

$$t(F,G) := \frac{\hom(F,G)}{|V(G)|^{|V(F)|}}.$$

A sequence of graphs $G_1, G_2, \dots$ is said to *converge* if the sequence $t(F, G_1), t(F, G_2), \dots$ converges for every graph F. Convergent sequences of graphs were defined and studied in [1, 2]. The main result of Lovász and Szegedy [9] is that for every convergent sequence of graphs there is a limit object in the form of a *graphon*, which is a symmetric measurable function $W\colon [0,1]^2 \to [0,1]$ (here *symmetric* means that $W(x,y) = W(y,x)$) such that $t(F, G_n) \to t(F, W)$ as $n \to \infty$ for all graphs F. Here $t(F, W)$ is defined by

$$t(F,W) := \int_{[0,1]^{V(F)}} \prod_{ij \in E(F)} W(x_i, x_j)\; dx_1 dx_2 \cdots dx_{|V(F)|}$$

The natural extension of these limits to hypergraphs was considered by Elek and Szegedy [4]. They constructed using ultraproducts an "ultralimit hypergraph" for any sequence of hypergraphs, and established a correspondence principle which enabled them to convert statements about finite hypergraphs, such as hypergraph regularity and removal lemmas [6, 11, 12], to measure-theoretic claims about ultralimit spaces. One of the consequences of their work is the existence of a limit object in the form of a measurable functions $W\colon [0,1]^{2^k-2} \to [0,1]$ for any convergent sequence of k-uniform hypergraphs.

In this paper we provide an ultraproduct-free construction of hypergraph limits. Our approach follows the original approach of Lovász and Szegedy, with the key ingredients being a weak Frieze-Kannan [5] type regularity lemma and the Martingale Convergence Theorem.

Date: February 8, 2013.

Convergence and limit object. For any k-uniform hypergraphs F and H, let $\hom(F,H)$ denote the number of homomorphisms from F to H, i.e., maps $V(F) \to V(H)$ that carry every edge of F to an edge of H. Define $t(F,H) := \hom(F,G)/\left|V(G)\right|^{|V(F)|}$. This is the probability that a random map $V(F) \to V(H)$ is a homomorphism. We say that a sequence of k-uniform hypergraphs $H_1, H_2, \dots$ *converges* if $t(F,H_1), t(F,H_2), \dots$ converges for every k-uniform hypergraph F.

For any positive integer n, define $[n] := \{1, 2, \dots, n\}$. For any set A, define $r(A)$ to be the collection of all nonempty subsets of A, and $r_<(A)$ to be collection of all nonempty proper subsets of A. More generally, let $r(A,m)$ denote the collection of all nonempty subsets of A of size at most m. So for instance, $r_<([n]) = r([n], n-1)$.

Any permutation σ of a set A induces a permutation on $r(A,m)$. We say that a function $W\colon [0,1]^{r([k],m)} \to [0,1]$ is *symmetric* if it remains invariant under any permutation of the coordinates induced by any permutation of $[k]$. For example, $W\colon [0,1]^{r_<([3])} \to [0,1]$ being symmetric means that

$$W(x_1, x_2, x_3, x_{12}, x_{13}, x_{23}) = W(x_{\sigma 1}, x_{\sigma 2}, x_{\sigma 3}, x_{\sigma 1 \sigma 2}, x_{\sigma 1 \sigma 3}, x_{\sigma 2 \sigma 3}) \tag{1.1}$$

for any permutation σ of $[3]$. Here we used the shorthand x_i to mean $x_{\{i\}}$ and x_{ij} to mean $x_{\{i,j\}}$.

Definition 1.1. A *k-uniform hypergraphon* is a symmetric measurable function $W\colon [0,1]^{r_<([n])} \to [0,1]$.

Example 1.2. A 3-uniform hypergraphon is a measurable function $W\colon [0,1]^6 \to [0,1]$ satisfying the symmetry condition (1.1).

For any k-uniform hypergraph F and hypergraphon W, define the homomorphism density by

$$t(F,W) := \int_{[0,1]^{r(V(F),k-1)}} \prod_{A \in E(F)} W(\mathbf{x}_{r_<(A)}) \, d\mathbf{x}_{r(V(F),k-1)}$$

Our convention throughout the paper is that $\mathbf{x}_{\mathcal{A}} = (x_A : A \in \mathcal{A})$ is a vector whose coordinates are indexed by a collection $\mathcal{A}$ of sets. For any $\mathcal{B} \subseteq \mathcal{A}$ we write $\mathbf{x}_{\mathcal{B}} = (x_B : B \in \mathcal{B})$ to mean the restriction of the vector to the coordinates indexed by $\mathcal{B}$.

Example 1.3. If $K_4^{(3)}$ is the complete 3-uniform hypergraph on 4 vertices and W is a 3-uniform hypergraphon, then

$$\begin{aligned} t(K_4^{(3)}, W) = \int_{[0,1]^{10}} & W(x_1, x_2, x_3, x_{12}, x_{13}, x_{23}) W(x_1, x_2, x_4, x_{12}, x_{14}, x_{24}) W(x_1, x_3, x_4, x_{13}, x_{14}, x_{34}) \cdot \\ & \cdot W(x_2, x_3, x_4, x_{23}, x_{24}, x_{34}) \, dx_1 dx_2 dx_3 dx_4 dx_{12} dx_{13} dx_{14} dx_{23} dx_{24} dx_{34}. \end{aligned}$$

Every k-uniform hypergraph H can be represented as a k-uniform hypergraphon W^H as follows: divide $[0,1]$ into $|V(H)|$ equal-length intervals $\left\{I_1, I_2, \dots, I_{|V(H)|}\right\}$. Define

$$W^H(\mathbf{x}_{r_<([k])}) = \begin{cases} 1 & \text{if } x_{\{i\}} \in I_{a_i} \text{ for } i = 1, \dots, k \text{ and } \{a_1, a_2, \dots, a_k\} \text{ is an edge of } F, \\ 0 & \text{otherwise.} \end{cases}$$

In particular, W^H only depends on the first k coordinates. It can be alternatively described as taking the adjacency array of H and then adding $2^k - 2 - k$ extra free coordinates. It is not hard to see that $t(F,H) = t(F,W^H)$ for every k-uniform hypergraph F.

The main purpose of this paper is to give a new proof of the following result of Elek and Szegedy [4, Thm. 7] on the existence of hypergraph limits.

Theorem 1.4. *For every convergent sequence of k-uniform hypergraphs $H_1, H_2, \dots$, there exists a k-uniform hypergraphon W so that $t(F,H_n) \to t(F,W)$ as $n \to \infty$ for every k-uniform hypergraph F.*

It may initially seem somewhat strange that we need 6 coordinates to describe the limit of 3-uniform hypergraphs, whereas every 3-uniform hypergraph can be described in terms of a 3-dimensional adjacency array. These extra dimensions do not arise for limits of graphs, but they are essential for hypergraphs. Here is a standard example illustrating why functions of the form $[0,1]^3 \to [0,1]$ cannot capture the richness of 3-uniform hypergraph limits. Take $G_n \sim \mathbb{G}(n, 1/2)$ to be a sequence of graphs on n vertices, where each edge is generated with probability 1/2, and let H_n be the 3-uniform hypergraph whose edges are the triangles of G_n. Then with probability one, $t(F, H_n) \to 2^{-e_2(F)}$ for every 3-uniform hypergraph F, where $e_2(F)$ is the number of pairs of vertices of F that are contained in some edge of F. The limit of H_n is different from, say, the constant hypergraphon 1/2, which is the limit of a sequence of 3-uniform hypergraphs where every triple of vertices is taken to be an edge independently with probability 1/2. To describe the limit of H_n, we need to incorporate the limit of G_n into the data, and this is achieved by the three extra coordinates. Indeed, the limit of H_n is the 3-uniform hypergraphon with value 1 on $[0,1]^3 \times [0,1/2]^3$ and 0 elsewhere. Intuitively, while the first three coordinates encode the vertex types, the last three coordinates encode the vertex-pair types. This hypergraphon is $\{0,1\}$-valued since it is deterministic once the vertex and vertex-pairs types are set. If we modify the sequence H_n so that each triangle of G_n is included as an edge of H_n with some probability p independently, then the limiting hypergraphon would take the value p on $[0,1]^3 \times [0,1/2]^3$ and 0 elsewhere.

For k-uniform hypergraphs, we can similarly impose some structure at each level, corresponding to j-subsets of vertices, for every $1 \leq j \leq k$. This is why we need a coordinate for every proper subset of $[k]$ in the description of hypergraph limits.

Random hypergraph model. To further illustrate the involvement of the $2^k - 2$ coordinates in a hypergraphon, let us review the associated random hypergraph model.

Recall that if $W \colon [0,1]^2 \to [0,1]$ is a graphon, then we have the following natural random graph model $\mathbb{G}(n, W)$ on n vertices: choose i.i.d. uniform $x_1, x_2, \ldots, x_n \in [0,1]$, and let there be an edge between vertices i and j with probability $W(x_i, x_j)$ independently. It was shown [9, Cor. 2.6] using Azuma's inequality that $\mathbb{G}(n, W)$ converges to the limit W almost surely.

Similarly, a k-uniform hypergraphon W gives a natural model $\mathbb{G}(n, W)$ of a random k-uniform hypergraph on n vertices: choose a uniformly random $\mathbf{x}_{r([n],k-1)} \in [0,1]^{r([n],k-1)}$ and add the edge $B = \{i_1, \ldots, i_k\} \subseteq [n]$ with probability $W(\mathbf{x}_{r_<(B)})$ independently. Essentially the same proof for graphs carry over to show [4, Thm. 11] that $\mathbb{G}(n, W)$ converges to W as $n \to \infty$ with probability one. The random hypergraphs H_n from the previous paragraph arise from this model.

Analytic version and compactness. It will be convenient to prove an analytic version of Theorem 1.4. We say that a sequence of k-uniform hypergraphons $W_1, W_2, \ldots$ converges if $t(F, W_1), t(F, W_2), \ldots$ converges for every k-uniform hypergraph F.

Theorem 1.5. *For every convergent sequence of k-uniform hypergraphons $W_1, W_2, \ldots$, there exists a k-uniform hypergraphon W so that $t(F, W_n) \to t(F, W)$ as $n \to \infty$ for every k-uniform hypergraph F.*

We prove Theorem 1.5 by induction on k. In fact we prove a colored generalization of Theorem 1.5, stated in §3. The generalization is motivated by the need for a stronger induction hypothesis for proving Theorem 1.5. A similar notion for graphs was considered in [8].

There is a neat interpretation of Theorem 1.5 in terms of compactness, following Lovász and Szegedy [10]. Let $\mathcal{W}_0^{(k)}$ denote the set of k-uniform hypergraphons. Give $\mathcal{W}_0^{(k)}$ be the weakest topology for which the functions $t(F, \cdot)$ are continuous for every k-uniform hypergraph F. Identify W with W' if $t(F, W) = t(F, W')$ for every k-uniform hypergraph F. Call this topology the *left-convergence topology* of $\mathcal{W}_0^{(k)}$.

Corollary 1.6. *The space $\mathcal{W}_0^{(k)}$ with the left-convergence topology is compact.*

Proof. The space is metrizable with the metric $\delta(W, W') = \sum_{i \geq 1} 2^{-i} |t(F_i, W) - t(F_i, W')|$ where $(F_i)_i$ is some enumeration of all isomorphism classes of k-uniform hypergraphs. We know that compactness is equivalent to sequential compactness in metric spaces. For any sequence $W_1, W_2, \ldots$, by sequential compactness of $[0,1]^{\mathbb{N}}$ we can find a subseqeuence $W_{n_1}, W_{n_2}, \ldots$ so that $t(F, W_{n_i})$ converges as $i \to \infty$ for every F. By Theorem 1.5 there is some limit $W \in \mathcal{W}_0^{(k)}$, thereby proving sequential compactness. ∎

When $k = 2$, Lovász and Szegedy [10] showed that $\mathcal{W}_0^{(2)}$ is compact under the cut metric topology, and Borgs, Chayes, Lovász, Sós, and Vesztergombi [1] showed that the cut metric topology is equivalent to the left-convergence topology. Lovász and Szegedy interpreted the compactness with respect to the cut metric as an analytic form of the regularity lemma, and they showed that the compactness of the space of graphons easily implies strong versions of the regularity lemma. Unfortunately, for $k \geq 3$, we do not know of a useful extension of the cut metric to hypergraphs (and there are some reasons to believe that such a natural metric might be too much to ask for), and this remains one of the main difficulties in working with convergence and limits of hypergraphs. An interesting open direction is to describe a natural and useful description of the distance between two hypergraphs that would give rise to the same topology on the space of hypergraphons as the one induced by homomorphism densities. The work of Elek and Szegedy [4] as well as the proof given in this paper offer some hints at the underlying topology of the space of hypergraphons, but it seems difficult at the moment to give a clean formulation of our understanding.

Organization. In §2 we review Lovász and Szegedy's proof and then sketch the idea of our extension. The details are in §3.

2. PROOF IDEA

The cut norm is crucial to the study of graph limits. For any symmetric measurable function $W\colon [0,1]^2 \to \mathbb{R}$, the cut norm is defined by

$$\|W\|_\square := \sup_{S,T \subseteq [0,1]} \left| \int_{S \times T} W(x,y)\; dxdy \right|,$$

where S and T range over all measurable subsets of $[0,1]$. We have the identity

$$\|W\|_\square = \sup_{u,v\colon [0,1] \to [0,1]} \left| \int u(x) W(x,y) v(y)\; dxdy \right| \tag{2.1}$$

where u and v range over all measurable functions $[0,1] \to [0,1]$.

Recall that a graphon is a symmetric measurable function $W\colon [0,1]^2 \to [0,1]$. For any measure preserving bijection $\phi\colon [0,1] \to [0,1]$ and any graphon W, define W^ϕ by $W^\phi(x,y) = W(\phi(x), \phi(y))$. We define the cut distance between graphons by

$$\delta_\square(U, W) = \inf_\phi \|U^\phi - W\|_\square,$$

where the infimum is taken over all measure preserving bijections $\phi\colon [0,1] \to [0,1]$. The cut distance can be defined for pairs of graphs by considering their associated graphons, although we will stick with the analytic language here.

The following *counting lemma* tells us that graphs that are close in cut distance share similar subgraph statistics. We review the proof as we will need an extension to hypergraphs.

Lemma 2.1 (Counting lemma). *For any graphons U and W and any graph F, we have*

$$|t(F, U) - t(F, W)| \leq |E(F)|\, \delta_\square(U, W).$$

Proof. It suffices to show that $|t(F,U) - t(F,W)| \le |E(F)|\, \|U - W\|_\square$ since $t(F,U) = t(F,U^\phi)$. Let $n = |V(F)|$ and $E(F) = \{i_1j_1, i_2j_2, \dots, i_mj_m\}$. Write as a telescoping sum

$$\begin{aligned} &t(F,U) - t(F,W) \\ &= \int_{[0,1]^n} \left(\prod_{e=1}^m U(x_{i_e}, x_{j_e}) - \prod_{e=1}^m W(x_{i_e}, x_{j_e}) \right) dx_1 \cdots dx_n \\ &= \sum_{f=1}^m \int_{[0,1]^n} \left(\prod_{e=1}^{f-1} U(x_{i_e}, x_{j_e}) \right) \left(U(x_{i_f}, x_{j_f}) - W(x_{i_f}, x_{j_f}) \right) \left(\prod_{e=f+1}^{m} W(x_{i_e}, x_{j_e}) \right) dx_1 \cdots dx_n. \end{aligned}$$

Every term in the final sum is bounded by $\|U - W\|_\square$. Indeed, for the f-th term, if we fix all variables except x_{i_f} and x_{j_f}, then we have an expression of the form $\int_{[0,1]^2} (U(x_{i_f}, x_{j_f}) - W(x_{i_f}, x_{j_f})) u(x_{i_f}) v(x_{i_f})\, dx_{i_f} dx_{j_f}$ for some functions u and v, and we can bound it using (2.1). ■

For any graphon W and any partition $\mathcal{P}$ of $[0,1]$ into a finite set of measurable subsets, let $W_\mathcal{P}$ be the step function graphon obtained by W by replacing its value at $(x,y) \in P_i \times P_j$, for any $P_i, P_j \in \mathcal{P}$, by the average of W over $P_i \times P_j$ (if either P_i or P_j has measure zero, then assign the value of $W_\mathcal{P}$ on $P_i \times P_j$ arbitrarily).

An important tool in the construction of graph limits is the following weak regularity lemma due to Frieze and Kannan [5] (see also [10, Lem 3.1]). It can be proved by a mean-squared-density increment argument.

Lemma 2.2 (Weak regularity lemma). *For every $\varepsilon > 0$ and every symmetric measurable function $W \colon [0,1]^2 \to [0,1]$, there is some partition $\mathcal{P}$ of $[0,1]$ into at most $2^{2/\varepsilon^2}$ parts such that $\|W - W_\mathcal{P}\|_\square \le \varepsilon$.*

Lovász and Szegedy [10] showed that with respect to the cut metric, after identifying graphons with cut distance zero, the space of all graphons is compact. Equivalently:

Theorem 2.3 (Lovász and Szegedy [10]). *Every sequence $W_1, W_2, \dots$ of graphons contains a subsequence converging to some graphon W in cut distance.*

Let us recall the idea of the proof of Theorem 2.3. Let $\varepsilon > 0$, and apply the weak regularity lemma to approximate every W_n by a step function with a bounded number of steps. By replacing each W_n by some $W_n^{\phi_n}$, we may assume that the partition $\mathcal{P}_n$ divides $[0,1]$ into intervals. We may now take a subsequence so that the lengths of the intervals converge, and the values of $(W_n)_{\mathcal{P}_n}$ inside the boxes induced by the partition converge, i.e., the value inside the (i,j)-th box of $(W_n)_{\mathcal{P}_n}$ converge to some value as $n \to \infty$ (may be different limits for different (i,j)). Then in this subsequence, $(W_n)_{\mathcal{P}_n}$ converges pointwise almost everywhere to some limit U_1, which is also a step function.

Now repeat the same procedure with a smaller $\varepsilon' < \varepsilon$. We obtain new partitions $\mathcal{P}'_n$ which are refinements of previous partitions. Call the resulting limit U_2. Note that steps of $(W_n)_{\mathcal{P}'_n}$ are refinements of the steps of $(W_n)_{\mathcal{P}_n}$, and the values of the latter can be obtained from by the former averaging over each step. Thus a similar relation holds for U_2 and U_1.

Now we repeat this procedure for a sequence of ε_k tending to zero. We obtain a sequence $U_1, U_2, \dots$ of step functions so that each U_s can be obtained from U_t, $t > s$, by average over each step. It follows that if (X,Y) is a uniform random point in $[0,1]^2$, then the sequence $(U_1(X,Y), U_2(X,Y), \dots)$ is a martingale. Since every U_s is bounded, the Martingale Convergence Theorem[1] (see [13, Thm. 11.5]) implies that the martingale converges with probability 1, and hence there is some $U \colon [0,1]^2 \to [0,1]$ which is the pointwise almost everywhere limit of the U_s's. One then checks that U is the desired limit.

---

[1] The Martingale Convergence Theorem says that every L^1-bounded martingale converges almost surely. Our martingales are actually bounded uniformly within $[0,1]$

In summary, the above proof consists of two main steps:

1. For each error tolerance ε, apply a weak regularity lemma to get a finite-dimensional approximation of each graphon, and then take a subsequential limit.
2. For a decreasing sequence of ε, these limits form a martingale, so the existence of the limit as $\varepsilon \to 0$ follows by Martingale Convergence Theorem.

Our construction of hypergraphons follows the same steps. For the first step, we need an extension of the cut norm.

Definition 2.4 (Higher cut norm)**.** For any symmetric measurable function $W\colon [0,1]^{r_<([k])} \to \mathbb{R}$, define the $(k-1)$-*cut norm* to be

$$\begin{aligned}\|W\|_{\square^{k-1}} &:= \sup_{\substack{S_1,\dots,S_k \subseteq [0,1]^{r([k-1])}\\ \text{symmetric measurable}}} \left| \int_{\bigcap_{i=1}^k \pi^{-1}_{r([k]\setminus\{i\})}(S_i)} W(\mathbf{x}_{r_<([k])})\, d\mathbf{x}_{r_<([k])} \right| \\ &= \sup_{\substack{u_1,\dots,u_k\colon [0,1]^{r([k-1])}\to[0,1]\\ \text{symmetric measurable}}} \left| \int_{[0,1]^{r_<([k])}} W(\mathbf{x}_{r_<([k])}) \prod_{i=1}^k u_i(\mathbf{x}_{r([k]\setminus\{i\})})\, d\mathbf{x}_{r_<([k])} \right|,\end{aligned}$$

where $\pi_{r([k]\setminus\{i\})}$ is the projection of $[0,1]^{r_<([k])}$ onto $[0,1]^{r([k]\setminus\{i\})}$.

The usual cut norm corresponds to the case $k = 2$.

Example 2.5. For any symmetric measurable function $W\colon [0,1]^{r_<([3])} \to \mathbb{R}$ the 2-cut norm $\|W\|_{\square^2}$ is defined by

$$\sup_{u_1,u_2,u_3} \left| \int_{[0,1]^6} W(x_1,x_2,x_3,x_{12},x_{13},x_{23}) u_1(x_2,x_3,x_{23}) u_2(x_1,x_3,x_{13}) u_3(x_1,x_2,x_{12}) dx_1 dx_2 dx_3 dx_{12} dx_{13} dx_{23} \right|$$

where u_1, u_2, u_3 vary over all symmetric measurable functions $[0,1]^{r([2])} \to [0,1]$.

For 3-uniform hypergraphs, this cut norm corresponds to edge-density deviations with respect to arbitrary vertex-pair subsets, as oppose to merely vertex subsets as was considered in the original Frieze-Kannan work [5]. Similar notions were similar by Chung [3] in extending Szemerédi's regularity lemma to hypergraphs.

This higher cut norm is only an intermediate step in the induction towards establishing the limits. It does not give the generalization of the cut distance asked for at the end of the introduction.

The counting lemma, Lemma 2.1, has a straightforward generalization with the same telescoping sum proof. A key observation is that every two edges of a k-uniform hypergraph intersect in at most $k-1$ vertices, so each term in the telescoping sum expansion of $t(F,U) - t(F,W)$ can be bounded by the $(k-1)$-cut norm of the difference $U - W$.

Lemma 2.6 (Counting lemma)**.** *For any k-uniform hypergraphons U and W and any k-uniform hypergraph F, we have*

$$|t(F,U) - t(F,W)| \le |E(F)|\, \|U - W\|_{\square^{k-1}}$$

Every symmetric partition $\mathcal{P}$ of $[0,1]^{r([k-1])}$ (here symmetric means that every part is symmetric under any permutation of coordinates induced by a permutation of $[k-1]$) induces a symmetric partition $\pi^*(\mathcal{P})$ of $[0,1]^{r_<([k])}$ which is the minimal partition refining all $\pi^{-1}_{r([k]\setminus\{i\})}(\mathcal{P})$ for $i = 1, \dots, k$. Here $\pi_{r([k]\setminus\{i\})}\colon [0,1]^{r_<([k])} \to [0,1]^{r([k]\setminus\{i\})}$ is the projection. For example, if $k = 2$ and $\mathcal{P}$ is a collection of intervals $\{I_i\}$, then $\pi^*(\mathcal{P})$ is a collection of boxes $\{I_i \times I_j\}$. In general, the cells of $\pi^*(\mathcal{P})$ have the form $\bigcap_{i=1}^k \pi^{-1}_{r([k]\setminus\{i\})}(S_i)$ where $S_1, \dots, S_k$ vary over all k-tuples of cells in the partition $\mathcal{P}$.

For any k-uniform hypergraphon W and symmetric partition $\mathcal{P}$ of $[0,1]^{r([k-1])}$, define $W_{\mathcal{P}}$ to be the k-uniform hypergraphon which is the step function obtained from W by replacing its value on each cell of $\pi^*(\mathcal{P})$ by its value on the cell. It can alternatively be defined as the conditional expectation $\mathbb{E}[W \mid \pi^*(\mathcal{P})]$. The following weak regularity lemma follows from a straightforward extension of the energy increment argument used to prove Lemma 2.2. We will give a proof later of a colored generalization (Lemma 3.4).

Lemma 2.7 (Weak regularity lemma)**.** *For every $\varepsilon > 0$ and every k-uniform hypergraphon W there exists a symmetric partition $\mathcal{P}$ of $[0,1]^{r([k-1])}$ into at most $2^{k/\varepsilon^2}$ parts so that $\|W - W_{\mathcal{P}}\|_{\square^{k-1}} \le \varepsilon$.*

Now we sketch the construction of the limit hypergraphon. Start with a sequence $W_1, W_2, \dots$ of k-uniform hypergraphons. For each ε, we get a partition $\mathcal{P}_n = \{P_{n,1}, P_{n,2}, \dots, P_{n,\ell}\}$ of $[0,1]^{r([k-1])}$ regularizing each W_n. Any symmetric measurable subset $P \subseteq [0,1]^{r([k-1])}$ can be associated to a $(k-1)$-hypergraphon W^P by integrating out the top coordinate:

$$W^P(\mathbf{x}_{r_<([k])}) = \int_0^1 1_P(\mathbf{x}_{r([k])}) dx_{[k]}. \tag{2.2}$$

Applying induction on k we can restrict to a subsequence so that $W^{P_{n,i}}$ converges to some $(k-1)$-uniform hypergraphon Y^i as $n \to \infty$ for each $1 \le i \le \ell$. Since $\mathcal{P}_n$ is a partition, $\sum_{i=1}^{\ell} W^{P_{n,i}} = 1$ for each i, so that $\sum_{i=1}^{n} Y^i = 1$ as well. This allows us to create a new partition $\mathcal{P}_\infty = \{P_{\infty,1}, \dots, P_{\infty,\ell}\}$ of $[0,1]^{r([k-1])}$ defined by

$$P_{\infty,i} = \left\{ \mathbf{x}_{r([k-1])} : \sum_{j=0}^{i-1} Y^j(\mathbf{x}_{r_<([k-1])}) \le x_{[k-1]} < \sum_{j=0}^{i} Y^j(\mathbf{x}_{r_<([k-1])}) \right\}, \quad i = 1, 2, \dots, \ell.$$

(For $i = \ell$ the $<$ should be changed to a $\le$) Visually we are partitioning the cube $[0,1]^{r([k-1])}$ by stacking up the hypergraphons $Y^1, \dots, Y^n$ along the $x_{[k-1]}$ coordinate[2]. The partition $\mathcal{P}_\infty$ serves as a limit of $\mathcal{P}_n$.

We can create a k-uniform hypergraphon U_1 which is a step function with steps $\pi^*(\mathcal{P}_\infty)$ and values being the limiting values (restricting to subsequences if necessary) on the corresponding steps in $(W_n)_{\mathcal{P}_n}$, just as in the Lovász-Szegedy argument.

We can repeat the same argument for a smaller $\varepsilon' < \varepsilon$ and obtain a refinement $\mathcal{P}'_\infty$ of $\mathcal{P}_\infty$, and a U_2 such that U_1 can be obtained from U_2 by averaging over the steps. Continuing, we get a martingale $U_1, U_2, \dots$ so that we can apply the Martingale Convergence Theorem to deduce the limit.

In order to carefully handle the simultaneous convergence of $W^{P_{n,i}}$ as $n \to \infty$ for all $1 \le i \le \ell$, we shall introduce in the next section hypergraphon-vectors and formulate a stronger induction hypothesis.

3. LIMITS OF HYPERGRAPHON-VECTORS

In this section we prove a strengthening of Theorem 1.5. The strengthening is motivated by the need to have a stronger induction hypothesis in order to simultaneously handle a vector of hypergraphons. Define a *k-uniform hypergraphon-vector* to be a list $\mathbf{W} = (W^1, W^2, \dots)$ of k-uniform hypergraphons (it may be indexed by any finite or countable set, but we use $\mathbb{N}$ as default).

We need a "colored" version of homomorphism densities where each edge of a hypergraph F is embedded in a different hypergraphon. This is analogous to counting embeddings of F in a multipartite hypergraph where each vertex of F is pre-assigned to a part. For every k-uniform hypergraph F, hypergraphon-vector $\mathbf{W}$, and map $\alpha\colon E(F) \to \mathbb{N}$, define

---

[2] This is the point in the proof where the extra coordinates (e.g., the last three coordinates in a 3-uniform hypergraphon) begin to play a role even if we start with a sequence of hypergraphs. More precisely, the limit Y^i may not be $\{0,1\}$-valued even if each $W^{P_{n,i}}$ is.

$$t_\alpha(F, \mathbf{W}) := \int_{[0,1]^{r(V(F),k-1)}} \prod_{A \in E(F)} W^{\alpha(A)}(\mathbf{x}_{r_<(A)})\, d\mathbf{x}_{r(V(F),k-1)}.$$

A counting lemma holds by a telescoping argument similar to Lemma 2.1.

Lemma 3.1 (Counting lemma)**.** *For any k-uniform hypergraphon-vectors $\mathbf{U}$ and $\mathbf{W}$, k-uniform hypergraph F and map $\alpha\colon E(F) \to \mathbb{N}$, if $\left\|U^i - W^i\right\|_{\square^{k-1}} \le \varepsilon$ for all i in the image of α, then*

$$|t_\alpha(F, \mathbf{U}) - t_\alpha(F, \mathbf{W})| \le |E(F)|\, \varepsilon.$$

Proof. Let $V = V(F)$ and $n = |V|$ and $E(F) = \{A_1, \dots, A_m\}$. Write as a telescoping sum

$$\begin{aligned}
&t_\alpha(F, \mathbf{U}) - t_\alpha(F, \mathbf{W}) \\
&= \int_{[0,1]^{r(V,k-1)}} \left(\prod_{e=1}^{m} U^{\alpha(A_e)}(\mathbf{x}_{r_<(A_e)}) - \prod_{e=1}^{m} W^{\alpha(A_e)}(\mathbf{x}_{r_<(A_e)}) \right) d\mathbf{x}_{r(V,k-1)} \\
&= \sum_{f=1}^{m} \int_{[0,1]^{r(V,k-1)}} \left(\prod_{e=1}^{f-1} U^{\alpha(A_e)}(\mathbf{x}_{r_<(A_e)}) \right) \left(U^{\alpha(A_f)}(\mathbf{x}_{r_<(A_f)}) - W^{\alpha(A_f)}(\mathbf{x}_{r_<(A_f)}) \right) \cdot \\
&\qquad\qquad \cdot \left(\prod_{e=f+1}^{m} W^{\alpha(A_e)}(\mathbf{x}_{r_<(A_e)}) \right) d\mathbf{x}_{r(V,k-1)}.
\end{aligned}$$

The f-th term is bounded by $\left\|U^{\alpha(A_f)} - W^{\alpha(A_f)}\right\|_{\square^{k-1}} \le \varepsilon$. Indeed, if we fix all variables other than $\mathbf{x}_{r_<(A_f)}$ then all the factors except for $U^{\alpha(A_f)}(\mathbf{x}_{r_<(A_f)}) - W^{\alpha(A_f)}(\mathbf{x}_{r_<(A_f)})$ has the form $u(\mathbf{x}_{r(B)})$ for some $B \subsetneq A_f$, where B is the intersection of two edges, and hence the integral can be bounded by the $(k-1)$-cut norm. ■

We need to maintain partition and refinement relations of the parts of a partition. After applying (2.2), these relations become equations of the form $\sum W^{P_j} = 1$ or W^{P_i}. This motivates the following definition. For a k-hypergraphon-vector $\mathbf{W}$, a *linear relation* is defined to be an equation of the form $\sum_{i=1}^{\ell} c_i W^i = b$ where for some integer ℓ and some real numbers $c_1, \dots, c_\ell, b$. The right-hand side denotes the constant hypergraphon with value b.

Definition 3.2. A sequence $\mathbf{W}_1, \mathbf{W}_2, \dots$ of k-uniform hypergraphon-vectors is said to converge to another k-uniform hypergraphon-vector $\mathbf{W}_\infty$ (denoted $\mathbf{W}_n \to \mathbf{W}_\infty$) if

(a) $t_\alpha(F, \mathbf{W}_n) \to t_\alpha(F, \mathbf{W}_\infty)$ as $n \to \infty$ for every k-uniform hypergraph F and $\alpha\colon E(F) \to \mathbb{N}$; and
(b) Every linear relation satisfied by all but finite many $\mathbf{W}_n$'s, $n \in \mathbb{N}$, is also satisfied by $\mathbf{W}_\infty$.

Here is the main result of this section.

Theorem 3.3. *Every sequence $\mathbf{W}_1, \mathbf{W}_2, \dots$ of k-uniform hypergraphon-vectors contains a subsequence that converges to some k-uniform hypergraphon-vector $\mathbf{W}_\infty$.*

Theorem 1.5 follows by taking each $\mathbf{W}_n$ to be a single k-uniform hypergraphon W_n (either by restricting the index set to one element, repeating W_n infinitely many times, or padding the rest of vector by 0). In Theorem 1.5, any subsequential limit is the limit of the entire sequence since all the homomorphism densities are assumed to converge.

Lemma 3.4 (Weak regularity lemma)**.** *Given k-uniform hypergraphons $W^1, W^2, \dots, W^m$ and a symmetric partition $\mathcal{P}$ of $[0,1]^{r([k-1])}$, there exists a symmetric partition $\mathcal{P}'$ refining $\mathcal{P}$ into at most $|\mathcal{P}|\, 2^{km/\varepsilon^2}$ parts such that $\left\|W^i - W^i_{\mathcal{P}'}\right\|_{\square^{k-1}} \le \varepsilon$ for each $1 \le i \le m$.*

Proof. We build the partition incrementally, starting with $\mathcal{P}$. At at a given stage suppose the partition is $\mathcal{Q}$. If $\left\|W^i - W^i_{\mathcal{Q}}\right\|_{\square^{k-1}} \leq \varepsilon$ for every i then we stop. Otherwise there is some i and symmetric measurable subsets $S_1, \dots, S_k \subseteq [0,1]^{r([k-1])}$ such that $|\mathbb{E}[1_S(W_i - W_{i,\mathcal{Q}})]| > \varepsilon$, where $S = \bigcap_{j=1}^k \pi^{-1}_{r([k]\setminus\{j\})}(S_j)$. Now let $\mathcal{Q}'$ be the minimal partition refining $\mathcal{Q}$ and $S_1, \dots, S_k$. We have

$$\left\|W^i_{\mathcal{Q}'}\right\|_2^2 = \left\|W^i_{\mathcal{Q}}\right\|_2^2 + \left\|W^i_{\mathcal{Q}'} - W^i_{\mathcal{Q}}\right\|_2^2 > \left\|W^i_{\mathcal{Q}}\right\|_2^2 + \varepsilon^2.$$

The equality follows by Pythagorean Theorem, and the inequality follows from $\left\|W^i_{\mathcal{Q}'} - W^i_{\mathcal{Q}}\right\|_2^2 \geq \left(\mathbb{E}[1_S(W^i_{\mathcal{Q}'} - W^i_{\mathcal{Q}})]\right)^2 > \varepsilon^2$. Since $\left\|W^i_{\mathcal{Q}}\right\|_2 \leq 1$ for every i, this increment can occur at most m/ε^2 times. Each time we introduced at most k new sets into the partition, thereby increasing the number of parts by a factor of at most 2^k. Thus the final partition has at most $|\mathcal{P}|\, 2^{km/\varepsilon^2}$ parts. ■

Remark. By allowing measure-zero parts, we can require that every part of $\mathcal{P}$ is refined into exactly $\lceil 2^{km/\varepsilon^2} \rceil$ parts.

Proof of Theorem 3.3. When $k = 1$, a k-uniform hypergraphon is a just a value in $[0,1]$, so the result follows trivially by sequential compactness of $[0,1]^{\mathbb{N}}$. Now we induct on k.

Let $\varepsilon_s = 1/s$. For every s and n, apply Lemma 3.4 to find a symmetric partition $\mathcal{P}_{n,s}$ of $[0,1]^{r([k-1])}$ so that

- $|\mathcal{P}_{n,s}| = \ell_s$ for some integer ℓ_s independent of n.
- $\mathcal{P}_{n,s+1}$ is a refinement of $\mathcal{P}_{n,s}$ where every part of $\mathcal{P}_{n,s}$ is refined into an equal number of parts;
- $\left\|W^m_n - W^m_{n,s}\right\|_{\square^{k-1}} \leq \varepsilon_s$ for all n and all $m \leq s$, where $W^m_{n,s} := (W^m_n)_{\mathcal{P}_{n,s}}$;

Let us arbitrarily order the parts of every partition while respecting refinement in the sense that if P comes before Q in $\mathcal{P}_{n,s}$ then every subset of P comes before every subset of Q in $\mathcal{P}_{n,s+1}$. Write $P_{n,s,j}$ for the j-th part of $\mathcal{P}_{n,s}$.

Consider $(k-1)$-uniform hypergraphons $Y^{s,j}_n := W^{P_{n,s,j}}$ as in (2.2). Collect the Y's into hypergraphon-vectors $\mathbf{Y}_n = (Y^{s,j}_n)_{s,j}$ indexed by $\{(s,j) \in \mathbb{N} \times \mathbb{N} : j \leq \ell_s\}$. These hypergraphon-vectors all satisfy the same linear relations arising from refinements, i.e.,

$$\sum_{j=1}^{\ell_1} Y^{1,j}_n = 1 \quad \text{and} \quad \sum_{j'=(j-1)\ell_{s+1}/\ell_s+1}^{j\ell_{s+1}/\ell_s} Y^{s+1,j'}_n = Y^{s,j}_n \quad \text{for } s \geq 1 \text{ and } 1 \leq j \leq \ell_s. \tag{3.1}$$

Using the induction hypothesis, we can replace $(\mathbf{W}_n)_{n\geq 1}$ by a subsequence so that (after relabeling the subsequence as $(\mathbf{W}_n)_{n\geq 1}$)

- There is some $(k-1)$-uniform hypergraphon-vector $\mathbf{Y}_\infty = (Y^{s,j}_\infty)_{s,j}$ so that $\mathbf{Y}_n$ converges to $\mathbf{Y}_\infty$.
- For every $m \leq s$ and $1 \leq j_1, \dots, j_k \leq \ell_s$, the common value $w_{n,m,s,j_1,\dots,j_k}$ of $W^m_{n,s}$ on $\bigcap_{i=1}^k \pi^{-1}_{r([k]\setminus\{i\})}(P_{n,s,j_i})$ converges to some value $w_{\infty,m,s,j_1,\dots,j_k}$ as $n \to \infty$.

For each s, let $\mathcal{P}_{\infty,s} = \{P_{\infty,s,1}, P_{\infty,s,2}, \dots, P_{\infty,s,\ell_s}\}$ be the symmetric partition of $[0,1]^{r([k-1])}$ defined by

$$P_{\infty,s,i} = \left\{ \mathbf{x}_{r([k-1])} : \sum_{j=0}^{i-1} Y^{s,j}_\infty(\mathbf{x}_{r_<([k-1])}) \leq x_{[k-1]} < \sum_{j=0}^{i} Y^{s,j}_\infty(\mathbf{x}_{r_<([k-1])}) \right\}, \quad i = 1, 2, \dots, \ell.$$

(when $i = \ell$ change the $<$ to a $\leq$) Visually we are forming a new partition of the cube $[0,1]^{r([k-1])}$ by stacking $Y^{s,1}_\infty, Y^{s,2}_\infty, \dots, Y^{s,\ell_s}_\infty$ along the $x_{[k-1]}$ coordinate direction.

Next, construct a k-uniform hypergraphon $W^m_{\infty,s} \colon [0,1]^{r_<([k])} \to [0,1]$ by setting it to have value $w_{\infty,m,s,j_1,\dots,j_k}$ on $\bigcap_{i=1}^k \pi^{-1}_{r([k]\setminus\{i\})}(P_{\infty,s,j_i})$. Construct k-uniform hypergraphon-vectors $\mathbf{W}_{n,s} =$

$(W_{n,s}^m)_m$ for $n \in \mathbb{N} \cup \{\infty\}$. By Lemma 3.5 we have for every s

$$t_\alpha(F, \mathbf{W}_{n,s}) \to t_\alpha(F, \mathbf{W}_{\infty,s}), \quad \text{as } n \to \infty. \tag{3.2}$$

For each s, let $\mathcal{B}_s$ denote the σ-algebra on $[0,1]^{r_<([k])}$ generated by $\pi^*(\mathcal{P}_{\infty,s})$ (see definition before Lemma 2.7) Then $\mathcal{B}_s \subseteq \mathcal{B}_{s+1}$ since $\mathcal{P}_{\infty,s+1}$ is a refinement of $\mathcal{P}_{\infty,s}$ for every s, as the linear relations (3.1) of $\mathbf{Y}_n$ corresponding to partitions and refinements are preserved in the limit.

Let X be a uniform random point of $[0,1]^{r_<([k])}$. Then $W_{\infty,s}^m$ is a $\mathcal{B}_s$-measurable random variable. Since $W_{n,s}^m$ can be obtained from $W_{n,s+1}^m$ by averaging over the steps, the same holds in the limit (we defer a more detailed justification until the end):

$$\mathbb{E}[W_{\infty,s+1}^m \mid \mathcal{B}_s] = W_{\infty,s}^m. \tag{3.3}$$

It follows that for each m, $(W_{\infty,1}^m, W_{\infty,2}^m, \dots)$ is a martingale relative to the filtration $(\mathcal{B}_1, \mathcal{B}_2, \dots)$. Since $W_{\infty,s}^m$ is bounded by 1 everywhere, it follows by Martingale Convergence Theorem that the martingale converges almost surely. Thus there is some hypergraphon W_∞^m for every $m \geq 1$ so that

$$W_{\infty,s}^m \to W_\infty^m \text{ pointwise almost everywhere as } s \to \infty. \tag{3.4}$$

Let $\mathbf{W}_\infty = (W_\infty^m)_m$. Any linear relation satisfied by $\mathbf{W}_n$ is preserved at each step of the operation and thus holds in the limit. For any k-uniform hypergraph F and $\alpha\colon E(F) \to \mathbb{N}$, we have

$$\begin{aligned}
&|t_\alpha(F, \mathbf{W}_n) - t_\alpha(F, \mathbf{W}_\infty)| \\
&\quad \leq |t_\alpha(F, \mathbf{W}_n) - t_\alpha(F, \mathbf{W}_{n,s})| + |t_\alpha(F, \mathbf{W}_{n,s}) - t_\alpha(F, \mathbf{W}_{\infty,s})| + |t_\alpha(F, \mathbf{W}_{\infty,s}) - t_\alpha(F, \mathbf{W}_\infty)| \\
&\quad \leq |E(F)|\, \varepsilon_s + |t_\alpha(F, \mathbf{W}_{n,s}) - t_\alpha(F, \mathbf{W}_{\infty,s})| + |t_\alpha(F, \mathbf{W}_{\infty,s}) - t_\alpha(F, \mathbf{W}_\infty)|\,.
\end{aligned}$$

Taking limsup as $n \to \infty$ and using (3.2) yields

$$\limsup_{n\to\infty} |t_\alpha(F, \mathbf{W}_n)| \leq |E(F)|\, \varepsilon_s + |t_\alpha(F, \mathbf{W}_{\infty,s}) - t_\alpha(F, \mathbf{W}_\infty)|\,.$$

Let $s \to \infty$ and it follows by (3.4) that $t_\alpha(F, \mathbf{W}_n) \to t_\alpha(F, \mathbf{W}_\infty)$ as $n \to \infty$. ■

Lemma 3.5. *Let $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_\infty$ each be an ordered symmetric partitions of $[0,1]^{r([k-1])}$ into ℓ parts and let $\mathbf{Y}_n = (Y_n^j = W^{P_{n,j}} : j = 1, \dots, \ell)$, $n \in \mathbb{N} \cup \{\infty\}$, be the corresponding $(k-1)$-uniform hypergraphon-vectors. Assume that $\mathbf{Y}_n$ converges to $\mathbf{Y}_\infty$. Let $\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_\infty$ be k-uniform hypergraphon-vectors so that every hypergraphon in $\mathbf{W}_n$ is a step function that is piecewise constant on cells of $\pi^*(\mathcal{P}_n)$ (see definition before Lemma 2.7). Furthermore, assume that for all m, as $n \to \infty$, the common values of W_n^m on each cell of $\pi^*(\mathcal{P}_n)$ converge to the common values of W_∞^m on the corresponding cells of $\pi^*(\mathcal{P}_\infty)$. Then for any k-uniform hypergraph F and any map $\alpha\colon E(F) \to \mathbb{N}$ we have $t_\alpha(F, \mathbf{W}_n) \to t_\alpha(F, \mathbf{W}_\infty)$ as $n \to \infty$.*

Proof. We are given that every W_n^m, $n \in \mathbb{N} \cup \{\infty\}$, can be written as

$$W_n^m(\mathbf{x}_{r_<([k])}) = \sum_{j_1,\dots,j_k=1}^{\ell} w_{n,m,j_1,\dots,j_k} \prod_{i=1}^{k} 1_{P_{n,j_i}}(\mathbf{x}_{r([k]\setminus\{i\})})\; d\mathbf{x}_{r_<([k])}. \tag{3.5}$$

where $w_{n,m,j_1,\dots,j_k}$ are the common values of W_n^m on the cells of $\pi^*(\mathcal{P}_n)$. When we expand

$$t_\alpha(F, \mathbf{W}_n) = \int_{[0,1]^{r(V(F),k-1)}} \prod_{A \in E(F)} W_n^{\alpha(A)}(\mathbf{x}_{r_<(A)})\; d\mathbf{x}_{r(V(F),k-1)}.$$

using (3.5) we can write $t_\alpha(F, \mathbf{W}_n)$ as a sum of terms of the form

$$w_{n,*} \int_{[0,1]^{r(V(F),k-1)}} \prod_{B \in E_{k-1}(F)} 1_{P_{n,j_B}}(\mathbf{x}_{r(B)})\ d\mathbf{x}_{r(V(F),k-1)}$$

$$= w_{n,*} \int_{[0,1]^{r(V(F),k-2)}} \prod_{B \in E_{k-1}(F)} Y_n^{j_B}(\mathbf{x}_{r_<(B)})\ d\mathbf{x}_{r(V(F),k-2)}. \quad (3.6)$$

where $E_{k-1}(F)$ denotes the set of all size $k-1$ subsets of $V(F)$ that is contained in some edge of F, and $w_{n,*}$ is the product of some combination of $w_{n,m,j_1,\dots,j_k}$'s. The second line comes from the first by integrating out all the variables x_B with $|B| = k-1$. A key point here is that the sets $P_{n,1}, \dots, P_{n,\ell}$ are disjoint, so that in every nonzero term in the expansion, each $B \in E_{k-1}(F)$ is associated to exactly one $1_{P_{n,j}}$, even if B is contained in several edges of F.

It remains to show that (3.6) converges to the correct limit as $n \to \infty$. We are given that $w_{n,m,j_1,\dots,j_k} \to w_{\infty,m,j_1,\dots,j_k}$ as $n \to \infty$. The remaining integral converges since $\mathbf{Y}_n \to \mathbf{Y}_\infty$. ■

Proof of (3.3). Fix s and $1 \le j_1, \dots, j_k \le \ell_s$, and for each $1 \le i \le k$ let J_i denote the set of indices so that $\{P_{n,s+1,j} : j \in J_i\}$ is a partition of P_{n,s,j_i} for every $n \in \mathbb{N} \cup \{\infty\}$. Then (3.3) amounts to the claim that

$$\int_{[0,1]^{r_<([k])}} \sum_{j'_1 \in J_1, \dots, j'_k \in J_k} w_{n,m,s+1,j'_1,\dots,j'_k} \prod_{i=1}^{k} 1_{P_{n,s+1,j'_i}}(\mathbf{x}_{r([k]\setminus\{i\})})\ d\mathbf{x}_{r_<([k])}$$

$$= \int_{[0,1]^{r_<([k])}} w_{n,m,s,j_1,\dots,j_k} \prod_{i=1}^{k} 1_{P_{n,s,j_i}}(\mathbf{x}_{r([k]\setminus\{i\})})\ d\mathbf{x}_{r_<([k])} \quad (3.7)$$

holds when $n = \infty$. We know that (3.7) holds for every finite n since $\mathcal{P}_{n,s+1}$ is a refinement of $\mathcal{P}_{n,s}$ and the $w_{w,n,s,j_1,\dots,j_k}$'s are just the average values of W_n^m on the cells of $\pi^*(\mathcal{P}_{n,s})$. It remains to show that both sides of (3.7) individually converge to the $n = \infty$ limit as $n \to \infty$.

By integrating out the variables $x_{[k]\setminus\{i\}}$, $1 \le i \le k$, the right-hand side of (3.7), for instance, becomes

$$\int_{[0,1]^{r([k],k-2)}} w_{n,m,s,j_1,\dots,j_k} \prod_{i=1}^{k} Y_n^{s,j_i}(\mathbf{x}_{r_<([k]\setminus\{i\})})\ d\mathbf{x}_{r([k],k-2)} = w_{n,m,s,j_1,\dots,j_k} t_\alpha(K_k^{(k-1)}, \mathbf{Y}_n)$$

where α sends the edge $[k] \setminus \{i\}$ to the index (s, j_i). The limit at $n = \infty$ follows from $\mathbf{Y}_n \to \mathbf{Y}_\infty$ and $w_{n,m,s,j_1,\dots,j_k} \to w_{\infty,m,s,j_1,\dots,j_k}$. The limit for the left-hand side of (3.7) is similar. ■

Acknowledgments

The author would like to thank Jacob Fox, László Lovász, Jennifer Chayes and Christian Borgs for helpful conversations.

References

[1] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008.

[2] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Ann. of Math. (2)*, 176(1):151–219, 2012.

[3] F. R. K. Chung. Regularity lemmas for hypergraphs and quasi-randomness. *Random Structures Algorithms*, 2(2):241–252, 1991.

[4] G. Elek and B. Szegedy. A measure-theoretic approach to the theory of dense hypergraphs. *Adv. Math.*, 231(3-4):1731–1772, 2012.

[5] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.

[6] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166(3):897–946, 2007.
[7] L. Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American Mathematical Society, 2012.
[8] L. Lovász and B. Szegedy. Limits of compact decorated graphs. arXiv:1010.5155.
[9] L. Lovász and B. Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006.
[10] L. Lovász and B. Szegedy. Szemerédi's lemma for the analyst. *Geom. Funct. Anal.*, 17(1):252–270, 2007.
[11] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006.
[12] T. Tao. A variant of the hypergraph removal lemma. *J. Combin. Theory Ser. A*, 113(7):1257–1280, 2006.
[13] D. Williams. *Probability with martingales.* Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge, 1991.

DEPARTMENT OF MATHEMATICS, MIT, CAMBRIDGE, MA 02139-4307.
E-mail address: `yufeiz@math.mit.edu`